



Federal - Commercial Spectrum Data:
*Understanding Information Exchange Needs, Issues,
and Approaches*

Anant Sahai
Kate Harrison

NITRD Wireless Spectrum R&D Senior Steering Group
Workshop VI Report

May 2015

Table of Contents

1	Executive Summary.....	4
2	Background	6
3	Workshop Synopsis and Participants.....	7
4	Research Recommendations	7
4.1	Understand the Needs and Tradeoffs.....	7
4.1.1	What Data is needed?.....	7
4.1.2	What location and time accuracy is required for protection from interference?	8
4.1.3	What information should be shared across multiple SAS?.....	8
4.1.4	What data can be reliably sensed and how long is this data useful?	9
4.1.5	How do we promote diverse uses of SAS data?	9
4.2	Build Trust	10
4.2.1	How can coordinated users coexist with uncoordinated devices?.....	10
4.2.2	How can a trusted interchange between stakeholders be established?.....	10
4.2.3	Could the SAS be used to calculate aggregate interference?	10
4.2.4	Can service level agreements (SLAs) provide the certainty needed for business?.....	11
4.3	Describe the Threats	11
4.3.1	What does data security and privacy mean in the digital age?	11
4.3.2	How do we assess the risks of sharing information with one or more SAS?	12
4.3.3	What are the threat models and attack scenarios that apply to SAS systems?	12
4.3.4	What would a secure SAS architecture look like?.....	12
4.4	Find Solutions.....	13
4.4.1	What are the tradeoffs when using data obfuscation?	13
4.4.2	Can security techniques from other domains work for spectrum sharing?	13
4.4.3	Can the inherent nature of wireless transmissions be used to help manage security?	13
4.4.4	How can we detect suspicious queries?	14
4.4.5	How can data retention policies be used to improve security?	14
4.5	Improve Enforcement	14
4.5.1	Can access to more data improve enforcement?.....	14
4.5.2	Can and should the SAS provide data for enforcement?.....	15

4.5.3	How can we achieve balance between prevention, deterrence, and remediation?.....	15
4.5.4	What forms of incentives or disincentives exist and how do they affect new entrants? ..	15
4.5.5	How do we design certification processes to promote trust and innovation?	16
4.5.6	How much of the enforcement process can and should be automated?	16
4.5.7	Can we design a data set to support enforcement research?	16
5	Summary	17
5.1	Appendix A: WSRD and WSRD Workshop Organization.....	19
5.2	Appendix B: Participants in the WSRD SSG Workshop VI.....	21
5.3	Appendix C: Agenda for the WSRD SSG Workshop VI	22

1 Executive Summary

As the Presidential Memorandum¹ established and the PCAST report² confirmed, both economic growth and the strategic needs of the United States depend increasingly on access to wireless spectrum. Concurrently, the traditional approach of completely clearing a band of spectrum in order to reallocate it to another service has become economically infeasible. Acknowledging this, The PCAST report identified the spectrum access system, or SAS, to be a key enabler for moving from our legacy spectrum assignment system to a new spectrum sharing ecosystem. Key to building an SAS based ecosystem is the need for accurate, accessible, and secure user data; reluctance of stakeholders to share their data has been identified as a major impediment. WSRD Workshop VI: *Federal - Commercial Spectrum Data: Understanding Information Exchange Needs, Issues, and Approaches* was held in October, 2014, to identify the issues surrounding information exchange as it pertains to spectrum sharing, and to recommend appropriate R&D to mitigate them.

The workshop made evident that the need for data goes beyond basic interference avoidance. Shared data is necessary to support the entire wireless ecosystem that contributes to economic growth and national security. Data is needed for planning, development, and to build and maintain trust in the system as a whole as well as to maximize spectrum utilization efficiency whenever possible³.

Because spectrum users are often involved in sensitive operations, privacy and security concerns are seen as major disincentives for sharing data. On the commercial side it could involve proprietary or customer information; on the Federal side it may involve military, intelligence, or other classified operations. Concerns fall into two major categories. First, issues arise simply because we have data being collected to answer queries for a broad range of users. The second category involves issues specific to the spectrum domain and is often related to military operations. Much of this data has a transitory value tied to both space and time that could open up opportunities for both vulnerabilities and defense mechanisms to emerge.

Despite progress in technologies and methods, appropriate incentives and reassurances are still needed to get commercial and Federal stakeholders to share the data necessary to make spectrum sharing, at least in the case of the SAS ecosystem, to work on a scale that is meaningful. The issue of incentives is of such importance that the next WSRD workshop will address it exclusively. In the meantime, the research recommendations contained within this report reflect the need to balance the needs for near-term profits in the commercial sector, the blue-sky vision and innovation of the academic sector, and the missions of the Federal Agencies.

¹ Presidential Memorandum: *Unleashing the Wireless Broadband Revolution*; <http://www.whitehouse.gov/the-press-office/presidential-memorandum-unleashing-wireless-broadband-revolution>; June 2010

² Report to the President: *Realizing the Full Potential of Government-Held Spectrum to Spur Economic Growth*; http://www.whitehouse.gov/sites/default/files/microsites/ostp/pcast_spectrum_report_final_july_20_2012.pdf

³ This will require the collection of not only static data (e.g. rules and regulations, primary user locations, terrain data, et.) but also dynamic data (primary and secondary usage information, operating parameters, etc.).

We propose finding federal-commercial data and information exchange solutions that will be adaptable over the long term yet be useful in promoting near term incremental change. If done correctly, we can have progressive solutions that adapt and improve as our understanding and technology evolve. This requires research into new technologies and institutional structures that build trust and foster continuous learning and improvement. The research portfolio required is therefore both broad and deep. The WSRD SSG would like to draw special attention to the following five key research areas:

- **Understand the Needs and Tradeoffs:** Understand how much data and what level of detail is needed for business planning, system operation, and enforcement taking into account the tradeoff between data fidelity, privacy, and security
- **Build Trust:** Establish a trusted interchange between the DoD, NTIA, and FCC to coordinate what information the government agencies regard as sensitive versus what can be made publicly available
- **Describe the Threats:** Identify the threat models and attack scenarios that are opened up by SAS systems
- **Identify Solutions:** Determine what existing data security, obfuscation⁴ and privacy practices can be adapted from other domains, and which scenarios will require solutions unique to spectrum sharing
- **Improve Enforcement:** Agree on the goals of enforcement, the types of enforcement needed, and how to quantify enforcement needs in various sharing scenarios using data

None of the above has an immediate off-the-shelf answer and each is critically important to making progress toward substantial Federal and commercial spectrum sharing.

⁴ Obfuscation refers to the hiding of information via alteration of the data. There are several common forms of obfuscation such as: Dithering answers (e.g. burying true “no” answers in a flurry of random “no” answers; Dithering inputs (e.g. perturbing true location values by a random amount); Binning (e.g. asking for an age category rather than exact age).

2 Background

Economic growth and the strategic needs of the United States depend increasingly on access to wireless spectrum. Concurrently, the traditional approach of clearing a band usage in order to reallocate it to another user has become increasingly complex and costly. Legacy wireless systems are often still within their useful lifespans, and while re-engineering them may be technically possible, it often involves a prohibitively expensive redesign of the entire system. Alternatively, legacy systems with a sparse usage pattern in both space and time, in principle, could share spectrum with new users.

The PCAST⁵ report mentioned the idea of a spectrum access system, or SAS, as a key enabler of spectrum sharing. Simply put, an SAS represents a new infrastructure that can interface with wireless spectrum systems and help manage access to underutilized spectrum while minimizing harmful interference. The benefits of the SAS accrue to both the commercial and Federal spectrum user with bi-directional sharing of spectrum. Both parties gain access to additional spectrum while reducing the need to obtain additional exclusive spectrum. In addition, it would allow the appropriate government agencies, to train at home with the same “spectrum agile” wireless devices and systems that they use when abroad.

Television Whitespace databases exemplify early simple SAS systems. When queried, they reply yes or no to whether a particular channel can be used at a particular time and location. Answers are based on information provided by the FCC and protected entity registrations. Future SAS systems will operate in more dynamic environments and may provide a much richer set of services: if the TV Whitespace Databases are “map providers”, then the envisioned SAS systems will be spectrum “air-traffic-controllers” or even spectrum “market-makers”. These advanced models are already being explored as part of the current regulatory work in the 3.5GHz band.

⁵ Ditto

3 Workshop Synopsis and Participants

Dynamic information sharing and management creates opportunities in areas including network and business intelligence, devices, applications, public safety operations, and security; but developing the next generation of spectrum management tools is complex. It requires consensus among stakeholders on several important issues such as: the purpose of collecting and sharing information, the type and minimal amount of data required, obfuscation of data as needed, security, and enforcement.

To provide a forum for this conversation, the Wireless Spectrum Research and Development Senior Steering Group (WSRD SSG) held a workshop, *Federal-Commercial Spectrum Data: Understanding Information Exchange Needs, Issues, and Approaches*, on October 21, 2014, at the National Science Foundation facility, in Arlington, VA. Information gathered from that workshop has been used to develop this report.

Workshop participants represented a cross section of experts from the government, industry and academia (a list of attendees can be found in Appendix B). They were asked to address the following goals for the workshop:

- Examine the technical and enforcement aspects of data management and information exchange between Federal and non-federal entities
- Investigate data security and privacy principles, technologies, and implementation techniques
- Discuss methods for building trust among stakeholders

4 Research Recommendations

The complexity of the inter-related issues discussed resulted in a collection of comments and examples that arranged themselves loosely into the key research areas described in the Executive Summary. The findings themselves are arranged (not in any particular order) around potential research questions that were developed and discussed throughout the day.

4.1 Understand the Needs and Tradeoffs

4.1.1 What Data is needed?

In the TV whitespaces, devices report their locations to the TV whitespace database (acting as a SAS in this band) in order to obtain permission to utilize available spectrum. In order to accurately give this permission, the TV whitespace database must know when and where the incumbents are operating as well as their respective protected zones. However, the spectrum community has come to a consensus that some amount of data is needed for real-time interference protection for incumbents or primary users. In bands other than the TV whitespaces, this data requirement has yet to be determined.

While the SAS is expected to be operating in real time, not all aspects of the spectrum sharing ecosystem operate at this time scale. For example, a company considering investment in the shared white spaces is interested in the expected value of that spectrum based on general temporal and spatial availability (e.g. available 10% vs 90% of the time; in major cities vs. rural areas only). In contrast, system designers need to know in great detail the temporal and spatial availability, interference protection

criteria, methods, and primary characteristics in order to build robust protocols and systems. The type and amount of data needed to satisfy these two examples vary widely, but are still distinct from the white space example where data is needed for real-time interference protection for an incumbent. In these cases, research is needed to establish if aggregate statistics suffice. Can these statistics be inferred (e.g. via spectrum occupancy measurements) rather than explicitly stated? Which parties or systems will provide this data? Can an SAS help provide long-term usage statistics to help business planning?

In addition to knowing the entities that are using the spectrum, SAS designers also need to be aware of the characteristics of the band, its users, and the sharing mechanisms. For example, how often must data about the protected entities be updated? How frequently will each device be contacting the SAS? How quickly must the SAS respond to a request? The answers to these questions rely on the characteristics and needs of all systems in the band.

4.1.2 What location and time accuracy is required for protection from interference?

In order to protect an incumbent's system, it is important that no interfering transmissions occur within a certain range of the incumbent's receiver(s) when the receiver is active. The relevant range varies with the sensitivity of the receiver as well as the transmit power of the potential interferer. It also varies with the number of potential interferers and the nature of the local wireless propagation environment.

In practice, we tend to overestimate this range by a certain margin to ensure interference protection for the primary user. But what if we cannot know exactly where (or when) a primary user will be active? In that case, additional buffers to this range (in location and time) can give the primary user the assurance that it will be protected at the cost of reducing spectrum availability for secondary users. This general methodology applies equally well to users with non-primary (but priority) access to the band.

Even when we could know the primary's usage pattern precisely, it may not always be desirable to reveal all aspects of the pattern⁶, e.g. in military applications. Consequently, it is important to understand the tradeoff between location accuracy, time accuracy, and the potential for efficient spectrum use. What does this tradeoff look like? What metrics can we use to assess this tradeoff?

How can we characterize the agility of primary or prioritized systems? What data constraints are required by incumbents or coordinated users for key functions such as clearing a channel? How does the agility of a primary or prioritized user affect users with a lower priority? How can we determine the suitability of sharing for a set of use cases?

4.1.3 What information should be shared across multiple SAS?

Today, TV white space SASs need to share information about protected entity registrations, e.g. wireless microphone users. These registrations are typically made in advance of the spectrum reservation which allows for lower latency in the synchronization of the multiple SAS.

⁶ This also has important primary user operational security and privacy implications. It is likely that some form of database obfuscation techniques will be needed to protect the operational security and privacy of primary users. Research has shown that there is a direct tradeoff between spectrum efficiency and level of operational security and privacy.

However, faster-moving incumbents, such as airborne radar systems, may not be able to give as much advance notice of their spectrum usage requirements. In this example, exchange of information among SASs may become difficult to do within the primary user's constraints.

The question becomes what types of information are needed to protect the incumbent? When is it practical to share information between SASs, how much data actually needs to be shared, and in some cases does it make sense to have a single SAS in order to meet these constraints?

4.1.4 What data can be reliably sensed and how long is this data useful?

The previous WSRD workshop, [*Understanding the Spectrum Environment: Using data and monitoring to improve spectrum utilization*](#), envisioned that there would be many SASs integrating spectrum sensing data in some way. This data may come from end-user sensor reports, reports from dedicated sensing devices, or some combination of the two.

Research is needed to know what the limits are for data that can be sensed by either an individual sensor or many collaborating sensors. How can we verify the reliability of sensing methods and the resulting data? Any particular sensing measurement describes a specific location at a point in time. As the measurement ages, the data (slowly or quickly) becomes irrelevant. What does this time scale look like for various incumbent systems (including radar-type systems)? Are there uses for sensing data that would not require security of the data? For example, can it be used to establish long-term usage trends in the band⁷? Is raw data required for these tasks or is aggregate data sufficient?

4.1.5 How do we promote diverse uses of SAS data?

SAS data can be used to improve the ecosystem, provide possible obfuscation and promote innovation. However, competition is critical in any ecosystem as a force that helps drive innovation. While the easiest solution is to have all SASs return identical answers to a given query, doing so is not necessary for protecting an incumbent from interference. What are the possible benefits to having SASs provide different but still "safe" answers to the same query? What does "safe" mean in this context?

If SAS's give different answers to the same query, what are the security implications of doing so? Does this undermine the effectiveness of some obfuscation techniques? Or will it have the opposite effect and make database inference attacks more difficult? How can obfuscation be architected and implemented in a way that does not prevent a diverse SAS ecosystem?

The sheer amount of data held by the SAS represents an enormous opportunity for both beneficial and malicious uses. For example, a SAS may use its query history to "steer" future users into less-used spectrum. On the other hand, it may use this same data to manipulate stock or commodity prices for its subscribers. What are the benefits and drawbacks of alternative uses of SAS data? How can these uses be monitored and malicious behavior detected and prevented?

⁷ This may trigger user privacy concerns.

4.2 Build Trust

4.2.1 How can coordinated users coexist with uncoordinated devices?

In addition to sharing between incumbents and secondary users, generally there will be sharing among secondary users⁸. In many cases, there will also be sharing among different secondary users. One possibility is to have the SAS manage this in a strictly prioritized manner, e.g. as suggested in the PCAST report's three-tier SAS proposal, but other possibilities exist and need to be explored.

For example, LTE systems are explicitly coordinated, both in terms of frequency planning as well as network operations such as handoff between cell towers. Meanwhile, traditional users of unlicensed bands (e.g. WiFi) are not explicitly coordinated - coordination happens implicitly through spectrum sensing. The wireless industry is currently trying to design LTE-U in a way that can survive despite the existence of uncoordinated users while also being a fair user of the spectrum.

The problem becomes even more complex when there is an incumbent in the band whose usage pattern injects variability into the spectrum availability. Does the SAS have a role in promoting coexistence between coordinated and uncoordinated users? What information would it need from both sides for this purpose? How does this change if the coordinated user has priority over the uncoordinated user but not primary rights to the band? What about the reverse situation?

4.2.2 How can a trusted interchange between stakeholders be established?

Is there a way to coordinate trusted real-time interchange of information the military regards as sensitive versus that which is publicly releasable? The military's current model of information classification distinguishes between data which is considered to be sensitive vs. open-source. Sensitive information can only be shared with people who can demonstrate the appropriate clearance. In principle, open-source information should be publicly available but in some cases it is not. For example, based on the workshop discussions, for the 3.6 GHz and 5 GHz spectrum bands, the operational maneuvering of US military units inside the US is classified even though it's claimed to be publicly inferable. If this information that is needed for possible spectrum sharing is open source, keeping it classified only serves as an unnecessary road block; the formal processes to declassify data introduce delays that impede the ability to open up spectrum for sharing in a timely manner.

How can we promote discussions about the true sensitivity of data? Which parties should be involved in these discussions? Most importantly: what is a long-term solution to this problem?

4.2.3 Could the SAS be used to calculate aggregate interference?

A major concern in spectrum sharing is the effect of aggregate interference on the incumbent. While any one or two users may not cause problems, if there are hundreds of devices trying to share spectrum being occupied by an incumbent, then the aggregate may amount to significant interference.

⁸ The IEEE 802.19.1 standard specifies radio technology independent methods for coexistence among dissimilar or independently operated wireless networks in TV white spaces.
http://standards.ieee.org/news/2014/ieee_802_19_1.html

While not a perfect solution, if the geo-location data of all the users were made available, the SAS could be used to calculate the aggregate interference before granting permission for a new device to use that same spectrum. This could require the SAS to have significant computational power, and even result in a reduced number of sharers in any given space, but it does provide an initial solution that moves the SAS beyond being a database to being a computational engine. This approach would be a significant improvement over current implementations and could spur the development of fast and more accurate computational propagation models.

4.2.4 Can service level agreements (SLAs) provide the certainty needed for business?

With the release or inference of the data from an SAS, there is some associated certainty in the measurements and predictions based on that data. How much certainty is needed for business planning such as where to install new infrastructure? One way of providing this certainty is via SLAs. For example, protected entities in the band may agree to operate only in city A at most 10% of the time and in city B at most 50% of the time. What are the incentives for protected entities to provide SLAs? Can this be a source of revenue?

In networking, ISPs service level agreements are common and are expressed at the statistical level. What is the right way to formulate and express an SLA in the spectrum sharing context? How can the agreement be monitored and enforced and what role should the SAS play in this?

Even when levels of service are not explicitly agreed upon by the users of the band, de facto or expected SLAs will form based on historical usage characteristics. These SLAs may even be enforceable through political processes, such as in the well-known garage door opener case. What is the role of the SAS in adding credibility to SLAs?

4.3 Describe the Threats

4.3.1 What does data security and privacy mean in the digital age?

With the advent of search engines and widespread access to vast quantities of data, we must rethink our notion of information security. For example, can the classified locations of military ships be determined simply by examining local newspapers and the social media accounts of soldiers' family members as some individuals claim?

Given the amount and type of information available through "open" sources, are there any additional vulnerability enabled by SAS-provided information? Can inference techniques exploit correlation between SAS-provided data and open-source data to discover sensitive information like the precise location of military devices? Should information stored in (but not directly provided by) the SAS be viewed as publicly available?

Finally, what is the threat model of SAS database inference? For example, should we consider adversaries that have the ability to integrate SAS query responses with spectrum sensing results? Should we consider adversaries that can collude in large numbers?

4.3.2 How do we assess the risks of sharing information with one or more SAS?

How do we balance the need for privacy, security, economic value, innovation, and system performance using quantitative and qualitative metrics? What are the fundamental tradeoffs among these properties given various SAS designs and security postures for both federal and commercial systems?

For example, it may be desirable to have multiple SAS in order to promote innovation and competition. However, this requires the exchange of information between SASs and makes them vulnerable to attack. Are there different SAS architectures that could better balance this risk against the benefits of more open and diverse systems? For example, a commercial LTE operator's security posture may be that its spectrum usage pattern should be kept private. How could coexistence with such an operator be managed? Should there be two types of SASs, one Federal and one commercial, to divide processing of sensitive and non-sensitive information?

4.3.3 What are the threat models and attack scenarios that apply to SAS systems?

Security researchers often define threat models and attack scenarios in order to assess the security of a system. For the same reasons, these models and scenarios need to be defined in the spectrum sharing context and used to design systems which effectively fend off attacks.

Which features make the spectrum sharing problem unique? For example, unlike the Internet's Domain Name System (DNS), SASs could incorporate real-time information from sensing data and be engaged in obfuscation.

Once we have designed a SAS that is not vulnerable to attacks, how can we verify and certify its security? Since it is impossible to predict all possible attack scenarios, it is important to put in place operational processes that will allow for quick recovery if there is an adverse event. What are these processes and how do we design and implement them?

4.3.4 What would a secure SAS architecture look like?

To design a secure SAS architecture, we need to use the library of generic cyberattacks that cybersecurity researchers have compiled and use them to develop specific defenses. This starts by identifying the classes of data security problems based on the type of system, type of data, time frame of vulnerability, and recipient of the information. Understanding the vulnerabilities in any given SAS architecture allows the use of design and practices to reduce them.

Without such an understanding, we risk designing systems and architectures with major vulnerabilities that may or may not be easily retrofitted. It is also important to understand the different classes of attack and their potential effects. For example, one can imagine an attack which reveals only historical data (e.g. usage patterns). How useful is this historical data to potential attackers and how does the utility of information for attackers decay with time? In contrast, we can also imagine attacks which reveal the future state of the system, or ways to control the future state, and could lead to denial-of-service for secondary users or loss of data security for the incumbent's system. Hence, in terms of data security, is a distributed SAS architecture more secure than a centralized SAS architecture? Are there other architectures that improve security? What are the drawbacks?

Understanding how to address vulnerabilities is also important. Will it require regulation? Should it be addressed before any deployment of the SAS or the device? Can it be corrected in real-time or after the unlikely event occurs?

4.4 Find Solutions

4.4.1 What are the tradeoffs when using data obfuscation?

Obfuscation can be used to allay privacy concerns that would otherwise have prevented any exchange of data. For example, if obfuscation can be used to protect military unit locations, it could enable spectrum sharing in military bands. The military may have additional reasons for sharing this information directly, for example, by sharing the approximate location of radar systems, there will be no desire for industry to build commercial off-the-shelf hardware that can localize these radar systems.

Are there spectrum specific metrics for data obfuscation or do these metrics need to be developed? For example, how can we quantify implementation cost and computational cost? How can we quantify the effect of obfuscation on spectrum utilization efficiency? Can obfuscation be done in a single unified international system that is compatible with different tradeoffs?

4.4.2 Can security techniques from other domains work for spectrum sharing?

The set of possible obfuscation techniques is very rich but generally the focus is on a statistical context, for example hiding individual incomes while revealing an average income. But spectrum-sharing is both statistical and specific --- while the user doesn't care about why it can't use a particular channel, the user definitely does care about being able to use that channel in that particular place. What new techniques need to be developed to address the problems in this area and what are the architectural implications?

4.4.3 Can the inherent nature of wireless transmissions be used to help manage security?

There are problems and opportunities that are unique to the spectrum sharing context. For example, the responses to database queries lead to a physical change in the world, namely wireless transmissions. As a result, the behavior of the database can be indirectly observed and identified in the real world without having to inspect the database itself or relying on software integrity and testing procedures.

Given these characteristics, which aspects of them can be used in a beneficial manner? For example, can we automatically detect database implementation bugs via spectrum sensing to complement the use of traditional software-engineering techniques? Similarly, can we detect and quantify regulatory "bugs" in the same way to complement traditional regulatory processes?

Which existing database security practices can be adapted to the spectrum sharing scenario?⁹ What new techniques need to be developed to be able to better contain the vulnerability footprint of a SAS design?

⁹ For example, the [cryptdDB](#) work shows how to allow certain database queries to be executed and used without the data being fully interpretable by the database provider. Instead, only users with appropriate keys can

4.4.4 How can we detect suspicious queries?

While the SASs primary function is responding to queries from secondary users, can it perform other functions as well? In particular, can it be used to detect suspicious queries? For example, can it detect queries that are meant to determine the precise location of a military installation? If true, should this be done in real-time or after the fact? Should this be detected by a suspicious-query detector in the SAS, a regulator, or a combination of both? Can results from other fields, e.g. anomaly detection, be used in this context?

4.4.5 How can data retention policies be used to improve security?

Enforcement, audit and debugging are three examples of auxiliary uses of SAS-held data. Unlike normal SAS operations, these applications may require historical data. For example, identifying a malfunctioning secondary device with high confidence may require a sufficiently long history of anomalous behavior. Similarly, identifying a malfunctioning primary registration system may require looking at the registration history. Are there additional use cases that should be considered? How much data and what kinds of data are needed for each use case?

Despite these possible uses, retaining data for a longer period of time increases the chances that it will be involved in a security breach. For example, an attacker with access to query the history for all devices in the Washington DC area may be able to discern the pattern of arrivals and departures of military units. This suggests that the SAS should try to purge information as often as possible.

Should these limits on data retention extend to secondary devices as well? For example, should secondary devices be required to purge historical query or sensing data? If so, how would this be accomplished and tested?

4.5 Improve Enforcement

4.5.1 Can access to more data improve enforcement?

Enforcement refers to the institutional structures and actions that ensure that the SAS and all parties involved act in a manner consistent with the regulations. Two common classes of enforcement that provide incentives for correct device behavior are *ex ante* (pre-action) and *ex post* (post-action). *Ex ante* includes such things as certification procedures and exclusion zones, while *ex post* include fines and injunctions.

To design the most efficient and effective enforcement mechanisms, we must improve our understanding of the goals of regulators as well as users categorized as primary, prioritized, secondary and general users. We need to quantify the effectiveness of enforcement, understand its theoretical and practical limits, and decide what is needed in various spectrum sharing scenarios. Is there a tradeoff between enforcement granularity and spectrum usage flexibility?

formulate queries and interpret the answers, while the heavy-duty database computations can be done in a way that is ignorant of the contents. Consequently, compromising the database or its cloud host doesn't lead to a loss of secrecy.

4.5.2 Can and should the SAS provide data for enforcement?

What types and amounts of data are needed to make reasonable and justified enforcement decisions? Can the SAS provide some of this data? What is the tradeoff between releasing this data and user privacy? How can we prevent enforcement actions inadvertently leaking sensitive information or serving as an avenue for attack on a SAS?

To the extent that the SAS has a role in enforcement, can database query responses be shaped to aid in enforcement? In a limited way existing data servers do this via “blacklist” functionality. Can and should this role be extended? If the SAS does not have a role in enforcement, who should take on this role? Are regulators uniquely positioned or could a third party be used?

How would an enforcement entity interface with the SAS to obtain the data required for enforcement and issue enforcement-related commands? What time scale is needed, how can it be met, and how might the needs and expectations for enforcement evolve over time?

4.5.3 How can we achieve balance between prevention, deterrence, and remediation?

Traditional spectrum management largely relies on ex ante enforcement via certification procedures and exclusion zones. To the extent that ex post enforcement is invoked, it tends to involve lengthy legal proceedings. These traditional procedures were developed in an era of static allocations and largely non-interactive devices.

How can we adjust these procedures for a modern world of spectrum sharing, dynamic devices, and SASs? Does a SAS change the balance needed between ex ante and ex post enforcement and allow for new forms of enforcement?

Are there scenarios when traditional ex post enforcement is undesirable? For example, a military entity may be reluctant to report cases of interference for fear of an attacker learning an effective attack strategy. Are there modifications to traditional ex post techniques that are suitable in these situations?

How does the desired balance between ex ante and ex post enforcement change depending on the users of the band? For example, the military may strongly prefer ex ante enforcement while another user may prefer ex post enforcement. Similarly, established companies may prefer the well-known ex ante enforcement but agile new entrants may prefer ex post. How can the primary and secondary users in a band agree upon this balance? Can remediation be used in cases where ex ante and ex post enforcement techniques are either too expensive to implement, useless, or otherwise undesirable?

4.5.4 What forms of incentives or disincentives exist and how do they affect new entrants?

Financial incentives can have different effects on various players in the secondary spectrum market. For example, a particular fine may be ineffective for a well-established company but drive a small innovative one into bankruptcy. A balanced enforcement mechanism must include non-financial incentives as well.

When is it appropriate to use financial vs. non-financial incentives? What non-financial incentives are available, how effective are they, and in which scenarios? How do they combine with ex ante enforcement and other mechanisms?

How can we design SASs and devices to be flexible or “future proof” in the enforcement sense? How can we tell if the approaches we take are effective?

4.5.5 How do we design certification processes to promote trust and innovation?

Traditional certification procedures are developed in the context of specific systems and their interactions. The process typically involves bringing the operators together and having them come up with a solution. This approach can be effective at generating certain types of outcomes but is limited by who is represented and is often biased against disruptive new technologies. How can we better represent industries and technologies that are too new, small, or decentralized to participate in these conversations? Who are the stakeholders and what are their roles and responsibilities?

What is the impact of lengthy certification processes on spectrum sharing? For example, can “formal methods” be applied such that algorithms are “correct by design” and participants have the ability to self-certify?

4.5.6 How much of the enforcement process can and should be automated?

Traditional ex post enforcement involves in-field measurements, equipment testing, and lengthy legal proceedings. The addition of SASs presents an opportunity for automating part, if not all, of the enforcement process. Ex ante approaches are usually more amenable to automation and cheaper to deploy than ex post. Are there situations that would require both?

How can enforcement be automated? Would automation require additional data to be collected? What are the challenges and the benefits of automated enforcement as compared to traditional enforcement? For example, does the ability to enforce on a short time scale allow us to mitigate other risks and be more agile for new uses?

4.5.7 Can we design a data set to support enforcement research?

To support spectrum enforcement research, it is critical to have access to a set of standardized training data. In other fields, data sets like intrusion traces have proved immensely beneficial for benchmarking and discovery. Which types of data would be useful for establishing benchmarks or discovering trends and underlying phenomena? Specifically, what are the important properties of the data that need to be preserved to make it useful? Can sensitive data be sanitized and still retain these properties? How can we keep them current and useful for future spectrum sharing scenarios? How do we obtain them and how can we mitigate the risks of making them public?

Is it possible to create a data set that allows researchers to discover the parameter ranges that have the most impact without requiring explicit sharing of federal or commercial data? For example, instead of revealing the actual operational characteristics of sensitive radars, one could imagine a simulated data set that allows the operational characteristics to vary over a wide range. Then researchers could

discover vulnerabilities and new approaches that work for a variety of parameter ranges. Regulators, who do have access to the sensitive data, could choose the approach that is relevant without leaking sensitive information to the public.

Finally, is there a “red team” approach to creating these data sets? For example, can a data set which incorporates (hidden or apparent) evidence of device misbehavior be released for the purposes of testing enforcement algorithms? Who can create and release such data sets?

5 Summary

In order for our Nation’s wireless ecosystem to grow and expand, we need to find ways to facilitate improved data sharing between stakeholders. Data is needed for planning, development, building trust, and maximizing efficient spectrum utilization.¹⁰

Because spectrum users are often involved in sensitive operations, privacy and security concerns are seen as major disincentives for sharing data. On the commercial side it could involve proprietary or customer information; on the Federal side it may involve military, intelligence, or other classified operations. Much of this data has a transitory value tied to both space and time that could open up opportunities for both vulnerabilities and defense mechanisms to emerge.

The experts that participated in the WSRD SSG workshop propose finding progressive solutions that improve as our understanding and technology evolve. This requires research into new technologies and institutional structures that build trust and foster continuous learning and improvement. The workshop drew special attention to the following five key areas:

- **Understand the Needs and Tradeoffs:** Understand how much data and what level of detail is needed for business planning, system operation, and enforcement taking into account the tradeoff between data fidelity, privacy, and security
- **Build Trust:** Establish a trusted interchange between the DoD, NTIA, and FCC to coordinate what information the government agencies regard as sensitive versus what can be made publicly available
- **Describe the Threats:** Identify the threat models and attack scenarios that are opened up by SAS systems
- **Identify Solutions:** Determine what existing data security, obfuscation¹¹ and privacy practices can be adapted from other domains, and which scenarios will require solutions unique to spectrum sharing

¹⁰ This will require the collection of not only static data (e.g. rules and regulations, primary user locations, terrain data, et.) but also dynamic data (primary and secondary usage information, operating parameters, etc.).

¹¹ Obfuscation refers to the hiding of information via alteration of the data. There are several common forms of obfuscation such as: Dithering answers (e.g. burying true “no” answers in a flurry of random “no” answers; Dithering inputs (e.g. perturbing true location values by a random amount); Binning (e.g. asking for an age category rather than exact age).

- **Improve Enforcement:** Agree on the goals of enforcement, the types of enforcement needed, and how to quantify enforcement needs in various sharing scenarios using data

None of the above has an immediate off-the-shelf answer and each is important to make progress toward the type of Federal and commercial spectrum sharing that will grow our economy and protect our national security.

5.1 Appendix A: WSRD and WSRD Workshop Organization

The Wireless Spectrum Research and Development Senior Steering Group (WSRD SSG) was established in 2010 to assist the Secretary of Commerce in creating and implementing a plan to facilitate research, development, experimentation, and testing to explore innovative spectrum-sharing technologies. Such an effort was called for by the June 28, 2010, *Presidential Memorandum: Unleashing the Wireless Broadband Revolution* as part of the overall effort to improve access to broadband services. Some 16 agencies participate in the WSRD SSG, which is convened under the auspices of the Networking and Information Technology Research and Development program (NITRD) Program. Realizing that progress in this area will require the involvement of the private and academic sectors as well as the federal agencies, the WSRD group was asked to focus on how to bring together the various research communities to collaborate on solutions. Following are the workshops that have been conducted by WSRD SSG.

Table 1: Prior WSRD Workshops		
WSRD I	Boulder, CO	July 26, 2011
WSRD II	Berkeley, CA	January 17-18, 2012
WSRD III	Boulder, CO	July 24, 2012
WSRD IV	Cambridge, MA	April 23-24, 2013
WSRD V	Arlington, VA	March 31, 2014

The five earlier workshops brought together key individuals from industry, academia, and the public sector with WSRD members to discuss research projects underway or proposed, wireless test beds, and specific areas of research interest. The focus of the first three workshops was on technology based research. While technology is a key ingredient in promoting wireless broadband growth and innovation, ensuring timely commercialization of technologies, creating successful business models, and establishing spectrum sharing practices also will require addressing a host of business, legal, and policy issues. Therefore the next two workshops broadened the scope and focused on promoting economic efficiency and understanding the spectrum environment in Workshop IV and V respectively.

This, our sixth workshop, was hosted by the National Science Foundation and was held in Arlington, VA on October 21, 2014.

The following were members of the workshop planning committee:

Chairman: Rangam Subramanian, NTIA

Byron Barker, NTIA

Jeff Boksiner, Army

Ira Keltz, FCC

Eric Nelson, NTIA-ITS

Jerry Park, Virginia Tech

Anant Sahai, University of California, Berkeley

Workshop Facilitator: Dan Mintz, ESEM Consulting

Workshop Coordinator: Ms. Wendy Wigen, NITRD

5.2 Appendix B: Participants in the WSRD SSG Workshop VI

Alder, Larry, Google
Barker, Byron, NTIA
Boksiner, Jeff, Army
Chang, Shawn, Congressional Staff
Chapin, John, DARPA
de Vries, Pierre, Silicon Flatirons
Doyle, Linda, Trinity College Dublin
Fette, Bruce, IDA
Gibson, Mark, Comsearch
Gupta, Anoop, Microsoft
Gurney, Dave, Motorola
Johnson, Mark, Navy
Kamal, Sherin, SAIC Inc.
Lackpour, Alex, Lockheed
Marshall, Preston, Representing WIN Forum
McDonald, Howard, DISA
McHenry, Mark, Shared Spectrum Co.
Mintz, Daniel (Dan), ESEM Consulting LLC
Mody, Apurva, BAE
Moorefield, Fred, DoD/CIO
Neel, James, Cognitive Radio Technologies, LLC
Nelson, Eric, NTIA/ITS.M
Park, Jung-Min (Jerry), VT
Prowell, Stacy, ORNL
Rennier, Tony, DMI
Sahai, Anant, Berkeley
Sharkey, Steve, TMO
Soltyka, Tony, DoD
Stine, John, MITRE
Subramanian, Vijayarangam (Rangam), NTIA
Tandon, Neeti, ATT
Tarazi, Iyad, Federated Wireless
Taylor, Tom, DoD CIO
Trappe, Wade, Rutgers University

5.3 Appendix C: Agenda for the WSRD SSG Workshop VI

8:00 AM	Continental Breakfast
8:30 AM	Welcome and Introductions: <i>Byron Barker, NTIA; Rangam Subramanian, NTIA</i>
8:45 AM	Keynote: Shawn Chang, Chief Democratic Counsel, House Energy and Commerce Committee
9:00 AM	Keynote: Fred Moorefield, Director, Spectrum Policy and Programs, Department of Defense, Chief Information Office
9:20 AM	Introduction to the Workshop Sessions: <i>Dan Mintz</i>
9:25 AM	Background Issues with Data Sharing: <i>Mark Gibson, Comsearch</i>
9:30 AM	Session I: Collecting and Sharing Data: Purpose, Uses, and Issues: <i>Moderated by Pierre de Vries, Silicon Flatirons</i> <i>Presenters: Neeti Tandon, AT&T; Mark Johnson, Navy; Linda Doyle, Trinity College Dublin;</i>
10:45 AM	Break
11:00 AM	Session II: Data Sharing and Obfuscation: State-of-the-Art Technologies and Research Needs: <i>Moderated by John Chapin, DARPA</i> <i>Presenters: Jeff Boksiner, Army; Jerry Park, Virginia Tech; John Chapin</i>
12:30 PM	Lunch
1:30 PM	Session III: Security of Data Storage, Access, and Real-time Delivery: Issues, Framework, and Obstacles: <i>Moderated by Wade Trappe, Rutgers</i> <i>Presenters: Wade Trappe; Howard McDonald, DISA; Stacy Prowell, ORNL</i>
3:00 PM	Break
3:15 PM	Session IV: Enforcement, Building Trust and Win-Win Collaboration: Principles, Framework and Next Steps: <i>Moderated by Mark Gibson</i> <i>Presenters: Pierre de Vries; Tom Taylor, OSD; Steve Sharkey, T-Mobile</i>
4:45 PM	Break
5:00-5:30 PM	Session V: Summary and Conclusions: <i>Moderated by Dan Mintz; Anant Sahai, Berkeley; Rangam Subramanian, NTIA</i>